

uplink 230 to common bus 240. Conversely, when $q(t)$ decreases below any of the programmable thresholds $Q(1) - Q(n)$, then formatter 252 at that time generates and sends a variable-rate control code to switching fabric 202 instructing it to increase the rate at which it transmits data for storage in buffer 250. Thus, for example, if there is a burst of locally switched

5 traffic on bus 240 such that data buffer 250 cannot empty its contents quickly and data buffer 250 begins to fill up, formatter 252 may generate one or more variable-rate control codes instructing switching fabric 202 to lower its data transmit rate one or more times. In this fashion, the data transmit rate is adjusted to the maximum sustainable level. It is noted that in this context, maximum sustainable level does not necessarily equate to 100% of the full transmit rate.

10 When the local burst subsides, data can flow out of buffer 250 at a rate greater than the rate at which buffer 250 receives data, thereby lowering the quantity of data stored within data buffer 250. In response, formatter 252 may generate one or more variable-rate control codes instructing switching fabric 202 to increase its data transmit rate one or more times. In this fashion, the data transmit rate is again adjusted to the maximum sustainable level.

Thus, in contrast to the prior art, should line card 204 be unable to accept high data rates, the data transmit rate will quickly lower to a sustainable rate instead of continuously cycling the fabric with start/stop control codes. In this configuration, in contrast to the prior art, no guess need be made as to the initial rate at which data is transferred to line card 204. Rather, as buffer 250 begins to backup with data, formatter 252 will send one or more variable-rate control codes to the switching fabric 202 instructing it to lower its data transmit rate one or more times. As the

15 buffer 250 empties, formatter 252 will send one or more variable-rate control codes to switching fabric 202 instructing it to increase its data transmit rate one or more times.

The present invention is described above as being embodied in line a card. The present invention should not be limited thereto. The present invention may be embodied in any device

25 that has a buffer for temporarily storing data received from some transmitting device. For example, the present invention may be embodied in switching fabric 202. Although not shown in the figures, switching fabric 202 may include first, second, and third input FIFO buffers for receiving and storing data from line cards 204, 206, and 208, respectively. Switching fabric 202 may include one or more formatters configured to generate and send control codes, including

30 variable-rate control codes, to line cards 204 through 208. For purposes of explanation,

switching fabric 202 will be described as having one formatter capable of generating and sending first, second, and third variable-rate control codes to line cards 204 through 208, respectively, for variably controlling the rate at which line cards 204 through 208 transmit data to switching fabric 202 for storage in respective buffers thereof.

Like the variable-rate control codes generated by formatter 252, the formatter of switching fabric 202 may generate the first, second, and third variable-rate control codes after comparing the quantity of data stored in the buffers against programmable thresholds. For purposes of explanation, the first, second, and third buffers of switching fabric 202 generate $qa(t)$, $qb(t)$, and $qc(t)$ representing the quantity of data stored in the first, second, and third buffers, respectively, at time t . In this explanatory embodiment, the switching fabric formatter may compare $qa(t)$, $qb(t)$, and $qc(t)$ against programmable thresholds $Qa(1)$ - $Qa(n)$, $Qb(1)$ - $Qb(n)$, and $Qc(1)$ - $Qc(n)$, respectively. When $qa(t)$, $qb(t)$, or $qc(t)$ increases beyond any of the programmable thresholds $Qa(1)$ - $Qa(n)$, $Qb(1)$ - $Qb(n)$, or $Qc(1)$ - $Qc(n)$, respectively, then at that time formatter 252 generates and sends the first, second, or third variable-rate control code to line cards 204, 206, or 208 respectively, instructing it to slow the rate at which it transmits data for storage in the first, second, or third buffer, respectively, of switching fabric 202. Conversely, when $qa(t)$, $qb(t)$, or $qc(t)$ decreases below any of the programmable thresholds $Qa(1)$ - $Qa(n)$, $Qb(1)$ - $Qb(n)$, or $Qc(1)$ - $Qc(n)$, respectively, then at that time formatter 252 generates and sends the first, second, or third variable-rate control code to line cards 204, 206, or 208 respectively, instructing it to increase the rate at which it transmits data for storage in the first, second, or third buffer, respectively, of switching fabric 202. Since switching fabric 202 is coupled to line cards 204 through 208 via downlinks 228, 232, and 236, respectively, two or more of the first, second, and third variable-rate control codes may be transmitted simultaneously.

Fabric 202 as shown in FIG. 2a may support two or more types of data transmission to the line cards. For example, switching fabric 202 may support high and low priority data traffic in a non-blocking manner where high priority traffic can pass low priority traffic. High priority traffic passing low priority traffic should not be limited to discrete packets. For example, suppose end device 222 (FIG. 2a) is sending low priority traffic to end device 210 via fabric 202 when end device 216 begins sending high priority traffic to end device 210 via fabric 202. Initially, fabric 202 chooses to send packets from line card 208 since no high priority traffic

is being sent to end device 210. But, when the high priority traffic appears from end device 216, fabric 202 can interrupt a low priority packet sent from line card 222 and send the entire high priority packet from end device 216, then resume transmission of the low priority packet.

FIG. 2c shows a line card 204 which may find use in the system 200 shown in FIG. 2a.

5 In contrast to the line card 204 shown in FIG. 2b, line card 204 shown in FIG. 2c includes two data buffers 260 and 262 which may be used to store distinct types of data transmitted from switching fabric 202. More particularly, data buffer 260 may be used to store data transmitted from switching fabric 202 having the highest priority, while data buffer 262 may store data transmitted from switching fabric 202 having the lowest priority. It is noted that the line card
10 shown in FIG. 2c may include further receiving buffers for storing data having priority which lies between the highest and lowest priorities mentioned above.

Line card 204 shown in FIG. 2c includes a formatter 268 coupled to a control circuit 270 and multiplexer 272. Formatter 268 is configured to generate control codes, including high and low priority variable-rate control codes, for controlling operational aspects of the switching fabric 202. The high and low priority variable-rate control codes are used to control the rate at which high and low priority data is transmitted to the first data buffer 260 and second data buffer 262, respectively. Multiplexer 272, in response to control circuit 270, couples the output of
15 formatter 268 to uplink 230 when formatter 268 generates one of its control codes. Otherwise, multiplexer 272 couples local bus 240 to uplink 230 so that data may be transmitted to switching fabric 202. Although not shown in FIG. 2c, line card 204 includes a switch for selectively
20 coupling the inputs of buffers 260 and 262 to downlink 228 depending on whether data received by the switch is designated high priority or low priority.

In FIG. 2c, formatter 268 is configured to receive information relating to the quantity of data stored within the buffers 260 and 262. For purposes of explanation formatter 268 will be
25 described as receiving $q1(t)$ and $q2(t)$ from buffers 260 and 262, respectively, where $q1(t)$ and $q2(t)$ represent the quantity of data stored in buffers 260 and 262, respectively, at time t , it being understood that formatter may be configured to generate $q1(t)$ and $q2(t)$ based on information provided by buffers 26a and 262. In response to receiving $q1(t)$ and $q2(t)$, formatter 268